

On Bitcoin and Red Balloons

Moshe Babaioff* Shahar Dobzinski† Sigal Oren ‡ Aviv Zohar§

November 14, 2011

Abstract

We study scenarios in which the goal is to ensure that some information will propagate through a large network of nodes. In these scenarios all nodes that are aware of the information compete for the same prize, and thus have an incentive *not* to propagate information. One example for such a scenario is the 2009 DARPA Network Challenge (finding red balloons). We give special attention to a second domain, Bitcoin, a decentralized electronic currency system.

Bitcoin, which has been getting a large amount of public attention over the last year, represents a radical new approach to monetary systems which has appeared in policy discussions and in the popular press [3, 11]. Its cryptographic fundamentals have largely held up even as its usage has become increasingly widespread. We find, however, that it exhibits a fundamental problem of a different nature, based on how its incentives are structured. We propose a modification to the protocol that can fix this problem.

Bitcoin relies on a peer-to-peer network to track transactions that are performed with the currency. For this purpose, every transaction a node learns about should be transmitted to its neighbors in the network. As the protocol is currently defined and implemented, it does not provide an incentive for nodes to broadcast transactions they are aware of. In fact, it provides a strong incentive not to do so. Our solution is to augment the protocol with a scheme that rewards information propagation. We show that our proposed scheme succeeds in setting the correct incentives, that it is Sybil-proof, and that it requires only a small payment overhead, all this is achieved with iterated elimination of dominated strategies. We provide lower bounds on the overhead that is required to implement schemes with the stronger solution concept of Dominant Strategies, indicating that such schemes might be impractical.

*Microsoft Research, Silicon Valley. moshe@microsoft.com

†Department of Computer Science, Cornell University. shahar@cs.cornell.edu

‡Department of Computer Science, Cornell University. sigal@cs.cornell.edu

§Microsoft Research, Silicon Valley. avivz@microsoft.com

1 Introduction

In 2009 DARPA announced the DARPA Network Challenge, in which participants competed to find ten red weather balloons that were dispersed across the United States [2]. Faced with the daunting task of locating balloons spread across a wide geographical area, participating teams attempted to recruit individuals from across the country to help. The winning team from MIT [9], incentivized balloon hunters to seek balloons by offering them rewards of \$2000 per balloon. Furthermore, after recognizing that notifying individuals from all over the US about these rewards is itself a difficult undertaking, the MIT team cleverly offered additional rewards of \$1000 to the person who directly recruited a balloon finder, a reward of \$500 to his recruiter, and so on. These additional payments created the incentive for participants to spread the word about MIT’s offer of rewards and were instrumental in the team’s success. In fact, the additional rewards are necessary: each additional balloon hunter competes with the participants in his vicinity, and reduces their chances of getting the prize.

MIT’s scheme still requires further improvement. As it is, a participant can create a fake identity, invite the fake identity to participate, and use that identity to recruit others. By doing so he increases his prize by 50%.¹ Thus the reward scheme should be carefully designed so it does not create an incentive for such a Sybil attack. Our goal is to design reward schemes that incentivize *information propagation* and counter the dis-incentive that arises from the competition from other nodes, and are *Sybil proof* (robust against Sybil attacks) while having a *low overhead* (a total reward that is not too high).

A related setting is a raffle, in which people purchase numbered tickets in hopes of winning some luxurious prize. Each ticket has the same probability of winning, and the prize is always allocated. As more tickets are sold, the probability that a certain ticket will win decreases. In this case again, there is a clear tension between the organizer of the raffle, who wants as many people to find out about the raffle, and the participants who have already purchased tickets and want to increase their individual chances of winning. The lesson here is simple, to make raffles more successful participants should be incentivized to spread the word. One example of a raffle already implementing this is Expedia’s “FriendTrips” in which the more friends you recruit the bigger your probability of winning.

As apparent from the previous examples, the tension between information propagation and an increased competition is a general problem. We identify an instantiation of this tension in the Bitcoin protocol, our main example for the rest of the paper. Bitcoin is a decentralized electronic currency system proposed by Satoshi Nakamoto in 2008 as an alternative to current government-backed currencies.² Bitcoin has been actively running since 2009. Its appeal lies in the ability to quickly transfer money over the internet, and in its relatively low transaction fees.³ As of November 2011, it has 7.5 million coins in circulation (called *Bitcoins*) which are traded at a value of approximately 3 USD per bitcoin. Below, we give a brief explanation of the Bitcoin protocol, and then explain where the incentives problem appears. We give a more comprehensive description of the protocol in Appendix B.

Bitcoin relies on a peer-to-peer network to verify and authorize all transactions that are performed with the currency. Suppose that Alice wants to reserve a hotel room for 30 bitcoins. Alice cryptographically signs a transaction to transfer 30 bitcoins from her to the hotel, and sends the signed transaction to a small number of nodes in the network. Each node in the network propagates the transaction to its neighbors. A node that receives the transaction verifies that Alice has signed it and that she does indeed own the bitcoins she is attempting to transfer. The node then tries to “authorize” the transaction⁴ by

¹Indeed, we have no evidence of such attacks in the DARPA challenge. If no such attacks were made, one possible explanation is the short time span of the challenge and its non-commercial, scientific essence. It seems quite plausible that if the challenge is repeated several times such attacks will become common.

²The real identity of Satoshi Nakamoto remains a mystery. A recent *New Yorker* article [3] attempts to identify him.

³There are additional properties that some consider as benefits: Bitcoins are not controlled by any government, and its supply will eventually be fixed. Additionally, it offers some degree of anonymity (although this fact has been contested [10]).

⁴In fact, a node authorizes a set of several transactions together. To keep the presentation simple, this simplified description considers only a single transaction. Our proposed solution applies to the case of multiple transactions as well.

attempting to solve a computationally hard problem (basically inverting a hash function). Once a node successfully authorizes a transaction, it sends the “proof” (the inverted hash) to all of its neighbors. They in turn, send the information to all of their neighbors and so on. Finally, all nodes in the network “agree” that Alice’s bitcoins have been transferred to the hotel.

In compensation for their efforts, nodes are offered a payment in bitcoins for successful authorizations. The system is currently in its initial stages, in which nodes are paid a predetermined amount of bitcoins that are created “out of thin air”. This also slowly builds up the bitcoins supply. But Bitcoin’s protocol specifies an exponentially decreasing rate of money creation that effectively sets a cap on the total number of bitcoins that will be in circulation. As this payment to nodes is slowly phased out, bitcoin owners that want their transactions approved are supposed to pay fees to the authorizing nodes.

This is where the incentive problem manifests itself. A node in the network has an incentive to keep the knowledge of any transaction that offers a fee for itself, as any other node that becomes aware of the transaction will compete to authorize the transaction first and claim the associated fee. For example, if only a single node is aware of a transaction, it can eliminate competition altogether by not distributing information further. With no competition, the informed node will eventually succeed in authorizing and collect the fee. The consequences of such behavior may be devastating: as only a single node in the network works to authorize each transaction, authorization is expected to take a very long time.⁵

We stress that false identities are a prominent concern in Bitcoin. In fact, the Bitcoin protocol is built around the assumption that nodes can create false identities, and considers a transaction fully approved only once nodes that control a majority of the CPU power in the network have accepted it, rather than just a majority of the nodes. The latter is vulnerable to Sybil attacks. Therefore any reward scheme for transaction distribution must discourage such attacks.

A Simplified Model

We present the model in the framework of Bitcoin. The presentation here is a bit informal, see Appendix A for a formal model. For simplicity assume that only one transaction needs to be authorized. We model the authorization process as divided into two phases: the first phase is a distribution phase. The second one is a computation phase, in which every node that has received the transaction is attempting to authorize it.

In general, it is common to think of an efficient peer to peer network as a random graph, in which messages propagate quickly, multiplying the number of recipients as the message is sent to each additional layer. Since we are not able to solve for the general case of random graphs, we will further simplify that and assume that the network consists of a forest of d -ary directed trees, each of them of height H .⁶ In the beginning of the *distribution phase* the buyer sends the details of the transaction to the t roots of the trees (which we shall term *seeds*). We think of the trees as directed, since the information (about the transaction) flows from the root towards the leaves. Let $n = t \cdot \frac{d^H - 1}{d - 1}$ be the total number of nodes.

In the distribution phase, each node v can send the transaction to any of its neighbors after adding any number of fake identities of itself before sending to that neighbor. All v ’s fake identities are connected to the same set of children. A node can condition its behavior only on the *length* of the chain above it, which can possibly include false identities that were produced by its ancestors.

In the *computation phase* each node that is aware of the transaction tries to authorize it. If there are k such nodes, each of them has the same probability of $\frac{1}{k}$ to authorize it first.

⁵Bitcoin’s difficulty level is adjusting automatically to account for the total amount of computation in the network. The expected time of a single machine to authorize a transaction is currently on the order of 60 days, instead of the 10 minutes it is expected to take if the entire network competes to authorize.

⁶The intuition for this simplification is that when d is small relatively to the number of nodes n , message propagation in a random graph resembles a tree. In some sense, the case of trees is harder than general graphs as each node monopolizes the flow of information to its descendants.

When a node p_h succeeds in authorizing a transaction we have a *winning chain* p_1, \dots, p_h , where p_1 is one of the seeds. Each node on the path in the tree from the seed to the authorizer appears on this chain (nodes are not able to remove their predecessors from the chain due to the use of cryptographic signatures), so the observed chain is a superset of the real path in the tree (it potentially includes duplicates). For allocating rewards on the “winning” chain we look at it in reverse order, reward r_1 is allocated to the authorizing node and r_h to the seed. We assume that there exists a minimal payment c that at least covers the expected computation cost for a single node to authorize the transaction by itself such that any node that is promised a reward of c will attempt to authorize the transaction. Thus, we require that the authorizer reward r_1 is at least c .⁷ For a smaller reward, nodes will refuse to participate in the first place (due to individual rationality). We normalize c to be 1.

Every reward scheme we consider in this paper is completely defined by non-negative rewards $r_{i,h}^s$ for every seed s and every $1 \leq i \leq h$. The reward $r_{i,h}^s$ is the reward given to the i -th identity on the winning chain of length h that starts at the authorizing node. A reward scheme is *individually rational* for \mathcal{H} in a tree rooted by s if for every $1 \leq h \leq \mathcal{H}$ we have that $r_{1,h}^s \geq 1$.

We want rewarding schemes that will incentivize *information propagation* and *no duplication*. That is, it will be in a node’s best interest to distribute the transaction to all its children without duplicating itself, as well as never duplicating when it authorizes. Our goal is to have information of the transaction reach the majority of the nodes in the network. We want to achieve this with small rewards, while minimizing the number of seeds (so the burden of the initial distribution is as low as possible).

Our Results

We start by introducing a new family of schemes: almost uniform schemes. Each member in the family is parameterized by a height parameter \mathcal{H} and a reward parameter β . Let v be the node that authorized the transaction and suppose that v is the l ’th node in the chain. If $l > \mathcal{H}$ no node is rewarded (so nodes “far” from the seed will not attempt to authorize the transaction). Otherwise, each node in the chain except v gets a reward of β , and v gets a reward of $1 + (\mathcal{H} - l + 1)\beta$. We show that if there are $\Omega(\beta^{-1})$ seeds, only strategy profiles that exhibit information propagation and no duplication survive every order of iterated removal of dominated strategies. Furthermore, there exists an order in which no other strategy profiles survive (see a discussion of this solution concept below).

This gives us two interesting schemes, for two different values of β . The first is when $\beta = 1$. In this case the $(1, \mathcal{H})$ -almost-uniform scheme requires only a constant number of seeds and the total payment is always $O(\mathcal{H})$. The second scheme is the $(\frac{1}{\mathcal{H}}, \mathcal{H})$ -almost-uniform scheme. This scheme works if the number of seeds is $\Omega(\mathcal{H})$. Its total payment is 2.

We combine both schemes to create a reward scheme that has *both* a constant number of seeds and a constant overhead. The *Hybrid Scheme* first distributes the transaction to a constant number of seeds using the $(1, 1 + \log_d H)$ -almost-uniform scheme. Then we can argue that at least H nodes are aware of the transaction. This fact enables us to further argue that the $(\frac{1}{H}, H)$ -almost uniform scheme guarantees that the transaction is distributed to trees of height H . At the end of the distribution phase, most of the nodes that are aware of the transaction receive a reward of $\frac{1}{H}$ if they are in the successful chain, so the expected overhead is low. We have the following (imprecisely stated) theorem:

Theorem: In the hybrid rewarding scheme, if the number of seeds $t \geq 14$, the only strategies that always survive iterated elimination of dominated strategies exhibit information propagation and no duplication. In addition, there exists an elimination order in which the only strategies that survive exhibit information propagation and no duplication. Furthermore, the expected total sum of payments is at most 3.

⁷In the actual Bitcoin protocol many transactions are authorized together, and thus the minimum reward r_1 per transaction is relatively small since the computation cost is split between many transactions.

Notice that this scheme exhibits in equilibrium low overhead, Sybil proofness, and provides the nodes with an incentive to propagate information.

Iterated removal of dominated strategies is the following common technique for solving games: first the set of surviving strategies of each player is initiated to all its strategies. Then at each step, a strategy of one of the players is eliminated from the set. The strategy that is eliminated is one that is dominated by some other strategy of the same player (with respect to the strategies in the surviving sets of all other players). This process continues until there is no strategy that can be removed. The solution concept prescribes that each player will only play some strategy from his surviving set. In general the surviving sets can depend on the order in which strategies are eliminated. Yet, we show that in our case, regardless of the order of elimination, profiles of strategies in which every node propagates information and never duplicates, survive. Moreover, for a specific order of elimination they are the only strategies that survive.

The intuition behind the elimination process is that decreasing competition by your own descendants might be unprofitable due to the distribution rewards. Consider one particular node. If the amount of external competition from non-descendent nodes is small, it prefers not to distribute the transaction, and thus increase the probability that it receives the reward for authorization. However, if sufficiently many non-descendent nodes are aware of the transaction, the node prefers to duplicate itself one less time, and thus distribute the transaction to its children and increase its potential distribution reward. Once all nodes increase distribution, the arms race begins: the competition each node faces is greater, and again it prefers to duplicate itself one less time and distribute all the way to its grand-children. As this process continues, all nodes eventually prefer to distribute fully and never to duplicate.

Iterated removal of dominated strategies is a strong solution concept, but ideally we would like our rewarding scheme to achieve all desired properties in the stronger notion of dominant strategies equilibrium. However, we show that in every dominant strategy scheme either the amount that the scheme must pay in equilibrium is huge, or the number of initial seeds t must be very large.

Theorem: Every individually rational reward scheme that propagates information to at least half of the network, and in which no-duplication and information-propagation are a dominant strategy for all nodes, has expected payment of at least $\frac{1}{10} \left(\frac{2^{H-4}}{t^2} + \frac{1}{t} \cdot \left(\frac{H-3}{t \cdot e} \right)^{H-3} \right)$, where t is the number of seeds.

Notice that for the sum of total rewards to be constant the number of seeds t has to be a significant part of the network. This implies that dominant strategy schemes are quite impractical.

Related Work

The paper describing Bitcoin’s principles was originally published as a white paper by Satoshi Nakamoto [8]. The protocol was since developed in an open-source project. To date, no formal document describes the protocol in its entirety, although the open source community maintains wiki pages devoted to that purpose [1]. Some research has been conducted with regards to its privacy [10]. Krugman discusses some related economic concerns [7].

The subject of incentives for information dissemination, especially in the context of social and peer-to-peer networks has received some attention in recent years. Kleinberg and Raghavan [6] consider incentivizing nodes to search for information in a social network. A node that possesses the information is rewarded for relaying it back, as are nodes along the path to it. A key difference between their model and ours, is that nodes either possess the sought-after information or do not. If they do, then there is no need for further propagation, and if they do not, they cannot themselves generate the information, and so do not compete with nodes they transmit to (they do however compete for the propagation rewards with other unrelated nodes).

Emek *et. al.* [5] consider reward schemes for social advertising scenarios. In their model, the goal of the scheme is to advertise to as many nodes as possible, while rewarding nodes for forwarding the

advertisement. Unlike their scenario, in our case the scheme is eventually only aware of a chain that results in a successful authorization, and only awards nodes on the path to the successful authorizer.

Other works consider propagation in social networks without the aspect of incentives. For example [4], which considers ways to detect propagation events and examines data from cellular call data.

Future Research

In this work we propose a novel low cost reward scheme that incentivizes information propagation and is Sybil proof. Currently we assume that the network is modeled as a forest of t d -ary trees. A challenging open question is to consider the setting where the network is modeled as a random d -regular graph. Another interesting extension to consider is one in which nodes have different computation power. That is, each node v has a computation power CPU_v , then the probability of a node v that is aware of the transaction to authorize it is $\frac{CPU_v}{\sum_u CPU_u}$. We leave the analysis of these models, as well as the implementation and empirical experimentation, to future research.

2 New Reward Schemes

Our main result is the Hybrid scheme, a reward scheme that requires only a constant number of seeds and a constant overhead. We will show that the only strategies that always survive iterated removal of dominated strategies are ones that exhibit information propagation and no duplication. The basic building blocks for this construction is a family of schemes, almost-uniform schemes, with less attractive properties. However, we will show that combining together two almost-uniform schemes enables us to obtain the improved properties of the Hybrid scheme.

2.1 The Basic Building Blocks: Almost-Uniform Schemes

The (β, \mathcal{H}) -almost-uniform scheme pays the authorizing node in a chain of length h a reward of $1 + \beta \cdot (\mathcal{H} - h + 1)$. The rest of the nodes in the chain get a reward of β . If the chain length is greater than \mathcal{H} no node receives a reward. The idea behind giving the authorizing node a reward of $1 + \beta \cdot (\mathcal{H} - h + 1)$ is to mimic the reward that the node would have gotten if it duplicated itself $\mathcal{H} - h$ times. Therefore, we can assume that no node duplicates itself before trying to authorize the transaction, and focus only on duplication before sending to its children.

Theorem 2.1 *Let $d \geq 3$. Suppose that the number of seeds is $t \geq 7$ and in addition there are initially at least $2\beta^{-1} + 6$ nodes except the t seeds that are aware of the transaction. Then, information propagation and no duplication are the only strategies that always survive iterated removal of weakly dominated strategies in the (β, \mathcal{H}) -almost-uniform scheme, for every node at depth \mathcal{H} . Furthermore, there is an elimination order in which these strategies are the only ones that survive. The total payment is $1 + \mathcal{H} \cdot \beta$.*

While the $2\beta^{-1} + 6$ additional nodes that are aware the transaction are not necessarily seeds, for now, one can think about them as additional seeds. The Hybrid scheme will exploit this extra flexibility to simultaneously obtain low overhead and small number of seeds. Before proving the theorem, let us mention two of its applications:

Corollary 2.2

1. *In the $(\frac{1}{\mathcal{H}}, \mathcal{H})$ -almost uniform scheme, if the number of seeds is at least $2\mathcal{H} + 13$, Then, information propagation and no duplication are the only strategies that always survive iterated removal of weakly dominated strategies in the (β, \mathcal{H}) -almost-uniform scheme. Furthermore, there is an elimination order in which these strategies are the only ones that survive. The total payment is 2.*

2. In the $(1, \mathcal{H})$ -almost uniform scheme, if the number of seeds is at least 15, Then, information propagation and no duplication are the only strategies that always survive iterated removal of weakly dominated strategies in the (β, \mathcal{H}) -almost-uniform scheme. Furthermore, there is an elimination order in which these strategies are the only ones that survive. The total payment is $\mathcal{H} + 1$.

Both parts of the corollary are direct applications of the theorem, by making sure that the number of initial seeds is at least $2\beta^{-1} + 13$. Later we introduce the Hybrid scheme that combines between two almost uniform schemes, one with $\beta = 1$ and one with $\beta = \frac{1}{\mathcal{H}}$, to obtain a scheme that uses both a constant number of seeds and its total expected payment is a small constant that does not depend on \mathcal{H} . The next subsection explains how to combine them together to obtain the Hybrid scheme.

The rest of the section is organized as follows. In Subsection 2.3 we prove our Theorem 2.1. The next subsection presents the Hybrid scheme. Finally, Subsection B.4 discusses how to implement the Hybrid scheme within the framework of the Bitcoin protocol.

2.2 The Final Construction: The Hybrid Scheme

We now present our main construction, the Hybrid scheme. The scheme works as follows: we run the $(\frac{1}{H}, H)$ -almost uniform scheme with a set of A seeds ($|A| = a$) and simultaneously run the $(1, 1 + \log H)$ -almost uniform scheme with a set of B seeds ($|B| = b$).

Theorem 2.3 *If it holds that $a \geq b \geq 7$, Then, information propagation and no duplication are the only strategies that always survive iterated removal of weakly dominated strategies in the (β, \mathcal{H}) -almost-uniform scheme (for every node of depth at most H in trees rooted by seeds in A , and $1 + \log H$ for every node rooted by seeds in B). Furthermore, there is an elimination order in which these strategies are the only ones that survive. The total payment is in expectation at most 3, and $2 + \log H$ in the worst case. A fraction of at least $\frac{a}{a+b}$ of the network is aware of the transaction after the distribution phase.*

Proof: The theorem is a consequence of Theorem 2.1: all nodes up to depth $\log H$ in trees rooted by nodes in B will propagate information and will not duplicate themselves. Thus, when considering the nodes in A , there are at least $7 \cdot \frac{d^{\log H} - 1}{d - 1} \geq 7H \geq 2H + 6$ nodes that are aware of the transaction, which are not part of the trees rooted by nodes in A . Thus we can apply again Theorem 2.1 to claim each seed in A roots a full tree of height H .

Notice that the worst case payment is $1 + \log H$ (because of the use of the $(1, 1 + \log H)$ -almost uniform scheme). The expected payment is $\frac{a \cdot \frac{d^H - 1}{d - 1} \cdot 2 + b \cdot \frac{dH - 1}{d - 1} \cdot (1 + \log H)}{a \cdot \frac{d^H - 1}{d - 1} + b \cdot \frac{dH - 1}{d - 1}} \leq 3$. As for the number of nodes that are aware of the transaction at the end of the distribution phase: all trees rooted by nodes in A are fully aware of the transaction. These nodes are a fraction of $\frac{a}{a+b}$ of the network (notice that we conservatively ignored nodes rooted by seeds in B that are aware of the transaction). \square

2.3 Proof of Theorem 2.1

We start with the following definition. The formal definition of the subgame uses the formal notation of Appendix A. The intuition is that in a φ -subgame a node that has $H - l$ ancestors (including clones), does not clone itself and propagates information if $l \leq \varphi + 1$. Otherwise for every child it duplicates itself at most $l - \varphi - 1$ times.

Definition 2.4 *The φ -subgame is the original game with restricted strategy spaces: the strategy space of node v includes only strategies such that for every remanning strategy profile of the rest of the players, every i and $l(v)$: $c_i^{l(v), v} \leq \max\{l(v) - \varphi - 1, 0\}$.*

The technical core of the proof is the following lemma:

Lemma 2.5 *In the φ -subgame, suppose that there are at least 7 seeds, and at least $2\beta^{-1} + 6$ additional nodes are aware of the transaction. Then, any strategy $(c_1^{l(v),v}, \dots, c_d^{l(v),v})$ of node v such that some $c_i^{l(v),v} = l(v) - \varphi - 1$ is dominated by the strategy $(c_1^{l(v)}, \dots, c_{i-1}^{l(v),v}, c_i^{l(v),v} - 1, c_{i+1}^{l(v),v}, \dots, c_d^{l(v),v})$.*

The proof of the lemma is in Subsection 2.3.1. Let us show how to use the lemma in order to prove the theorem. First, we observe that the 0-subgame is simply the original game. If the conditions of the lemma hold, we can apply the lemma to eliminate all strategies $(c_1^{l(v),v}, \dots, c_d^{l(v),v})$ such that some $c_i^{l(v),v} = l(v) - 1$. Notice that now we have a 1-subgame. Applying the lemma again we get a 2-subgame. Similarly, we repeatedly apply the lemma until we get a $(\mathcal{H} - 1)$ -subgame. Notice that the only surviving strategy of each node is $(0, \dots, 0)$ for every $l(v)$: that is each node propagates information and does not duplicate. This shows that there exists an elimination order in which every node propagates information and does not duplicate. Next, we prove that the strategy profile in which all nodes propagate information and do not duplicate survives *every* order of iterated elimination of dominated strategies process (proof in the appendix):

Lemma 2.6 *Let T be a sub-game that is reached via iterated elimination of dominated strategies in the (β, \mathcal{H}) -almost-uniform scheme, and suppose that there are at least 7 seeds, and at least $2\beta^{-1} + 6$ additional nodes are aware of the transaction. Then there exists a strategy profile $s \in T$ in which every node at depth at most \mathcal{H} fully propagates and do not duplicate.*

2.3.1 Proof of Lemma 2.5

Now, let us observe a node v and fix $l = l(v)$. The non-trivial case is when $l \geq \varphi + 1$. For convenience, we drop the superscript $(l(v), v)$ and denote a strategy by (c_1, \dots, c_d) . Without loss of generality we assume $c_1 \leq c_2 \leq \dots \leq c_d$. It will also be useful to define $y_i = l - c_i - 1$. Note that our previous assumption implies that $y_1 \geq y_2 \geq \dots \geq y_d$.

Let s_{-v} be a strategy profile of all other nodes except v , and denote by $A_{s_{-v}}(y_i)$ the size of the subtree rooted at the i 'th child of v that learns of the transaction. The utility of v is then:

$$U_v^{s_{-v}}(y_1, \dots, y_d) = \underbrace{\frac{1 + \beta \cdot l}{k + \sum_{i=1}^d A_{s_{-v}}(y_i)}}_{v \text{ authorizes}} + \underbrace{\frac{\sum_{i=1}^d \beta(l - y_i) \cdot A_{s_{-v}}(y_i)}{k + \sum_{i=1}^d A_{s_{-v}}(y_i)}}_{\text{a decedent of } v \text{ authorizes}}$$

where k is the number of nodes which are aware of the transaction outside the subtree which v is the root of (this number includes v itself). We show that (c_1, \dots, c_d) , where $c_d = l(v) - \varphi - 1$ is dominated by $(c_1, \dots, c_d - 1)$. For this purpose, it is sufficient to show the following claim:

Claim 2.7 *Let T_{-v} be the set of strategy profiles of all nodes other than v in which:*

1. *The subtree rooted at v 's d 'th child spans a complete d -ary tree of height y_d when v duplicates itself c_d times, or of a height $y_d + 1$ when v duplicates itself $c_d - 1$ times.*
2. *There are at least $k \geq 2\beta^{-1} + 6d\frac{d^y_d - 1}{d - 1} + 6$ nodes that are not descendants of v that are aware of the transaction.*

Then: $\forall s_{-v} \in T_{-v} \quad U_v^{s_{-v}}(c_1, \dots, c_d - 1) > U_v^{s_{-v}}(c_1, \dots, c_d)$

Notice, that in any φ -subgame, both of the conditions for the lemma hold in every profile: in a φ -subgame we have that if node v clones itself at most $l(v) - \varphi - 2$ times then his child will span a tree that contains a full d -ary tree of height φ , and we are guaranteed k nodes that are aware of the transaction.

We therefore turn to prove the claim.

Proof: The new strategy $(c_1, \dots, c_d - 1)$ in effect distributes the message to one additional layer of the subtree rooted at the d 'th child of v . That extra layers contains d^{y_d} nodes. The utility of this strategy is therefore:

$$U_v^{s-v}(y_1, \dots, y_{d-1}, y_d + 1) = \frac{1 + \beta \cdot l + \sum_{i=1}^{d-1} \beta(l - y_i) \cdot A_{s-v}(y_i) + \beta(l - y_d - 1)(A_{s-v}(y_d) + d^{y_d})}{k + \sum_{i=1}^d A_{s-v}(y_i) + d^{y_d}}$$

So we have to show that:

$$\forall s-v \quad U_v^{s-v}(y_1, \dots, y_{d-1}, y_d + 1) > U_v^{s-v}(y_1, \dots, y_d)$$

For convenience, we will drop the index $s-v$, but remember that we must check for every possible strategy profile of the other nodes that:

$$\frac{1 + \beta \cdot l + \sum_{i=1}^{d-1} \beta(l - y_i) \cdot A(y_i) + \beta(l - y_d - 1)(A(y_d) + d^{y_d})}{k + \sum_{i=1}^d A(y_i) + d^{y_d}} > \frac{1 + \beta \cdot l + \sum_{i=1}^d \beta(l - y_i) \cdot A(y_i)}{k + \sum_{i=1}^d A(y_i)}$$

multiplying by the denominator and dividing both sides by β :

$$\begin{aligned} \left(k + \sum_{i=1}^d A(y_i)\right) \cdot (l - y_d - 1)(A(y_d) + d^{y_d}) &> \left(k + \sum_{i=1}^d A(y_i)\right) (l - y_d) \cdot A(y_d) + d^{y_d} \cdot \left(\beta^{-1} + l + \sum_{i=1}^d (l - y_i) \cdot A(y_i)\right) \\ \left(k + \sum_{i=1}^d A(y_i)\right) \cdot (l - y_d - 1) \cdot d^{y_d} - \left(k + \sum_{i=1}^d A(y_i)\right) \cdot A(y_d) &> d^{y_d} \cdot \left(\beta^{-1} + l + \sum_{i=1}^d (l - y_i) \cdot A(y_i)\right) \\ \left(k + \sum_{i=1}^d A(y_i)\right) \cdot \left(l - y_d - 1 - \frac{A(y_d)}{d^{y_d}}\right) &> \left(\beta^{-1} + l + \sum_{i=1}^d (l - y_i) \cdot A(y_i)\right) \end{aligned}$$

Note that $l - y_d - 1 - \frac{A(y_d)}{d^{y_d}} > 0$ since $l - y_d \geq 2$ and $\frac{A(y_d)}{d^{y_d}} < 1$. Therefore we can divide by $l - y_d - 1 - \frac{A(y_d)}{d^{y_d}}$, after some rearranging we get that:

$$\begin{aligned} k &> \frac{\beta^{-1} + l + \sum_{i=1}^d (l - y_i) \cdot A(y_i)}{l - y_d - 1 - \frac{A(y_d)}{d^{y_d}}} - \sum_{i=1}^d A(y_i) \\ k &> \frac{\beta^{-1} + l + \sum_{i=1}^d (l - y_i - l + y_d + 1 + \frac{A(y_d)}{d^{y_d}}) \cdot A(y_i)}{l - y_d - 1 - \frac{A(y_d)}{d^{y_d}}} \\ k \cdot \left(l - y_d - 1 - \frac{A(y_d)}{d^{y_d}}\right) &> \beta^{-1} + l + \sum_{i=1}^d \left(-y_i + y_d + 1 + \frac{A(y_d)}{d^{y_d}}\right) \cdot A(y_i) = (*) \end{aligned}$$

Recall that that the d 's child spans a full tree of height y_d : $A(y_d) = \frac{d^{y_d} - 1}{d - 1}$. We bound the right hand side of Equation 2.3.1 as follows:

$$\begin{aligned} (*) &\leq \beta^{-1} + l + \sum_{i: y_i = y_d} \left(1 + \frac{A(y_d)}{d^{y_d}}\right) \cdot A(y_i) + \sum_{i: y_i = y_d + 1} \left(\frac{A(y_d)}{d^{y_d}}\right) \cdot A(y_i) \\ &\quad + \sum_{i: y_i > y_d + 1} \left(-1 + \frac{A(y_d)}{d^{y_d}}\right) \cdot A(y_i) \end{aligned}$$

The last summation is negative because $\frac{A(y_d)}{d^{y_d}} < 1$, and once we substitute $A(y_d)$ into the expression we have that:

$$(*) \leq \beta^{-1} + l + \sum_{i:y_i=y_d}^d \left(1 + \frac{A(y_d)}{d^{y_d}}\right) \cdot A(y_d) + \sum_{i:y_i=y_d+1}^d \frac{A(y_d)}{d^{y_d}} \cdot A(y_d + 1)$$

Now notice that:

$$\left(1 + \frac{A(y_d)}{d^{y_d}}\right) \cdot A(y_d) = \frac{A(y_d) + d^{y_d}}{d^{y_d}} \cdot A(y_d) = \frac{A(y_d + 1) \cdot A(y_d)}{d^{y_d}}$$

and so we can write:

$$(*) \leq \beta^{-1} + l + \sum_{i:y_i=y_d \vee y_i=y_d+1} \left(1 + \frac{A(y_d)}{d^{y_d}}\right) \cdot A(y_d) \leq \beta^{-1} + l + \sum_{i:y_i=y_d \vee y_i=y_d+1} 2 \cdot A(y_d) \quad (1)$$

$$\leq \beta^{-1} + l + 2 \cdot d \cdot A(y_d) \quad (2)$$

Therefore, we are looking for k such that:

$$k > \frac{\beta^{-1} + l + 2 \cdot d \cdot A(y_d)}{l - y_d - 1 - \frac{A(y_d)}{d^{y_d}}} = (**)$$

We now divide into two cases.

Case I: $l \geq 2y_d + 4$. In this case By using the fact $\frac{A(y_d)}{d^{y_d}} < 1$ and continuing from Equation (2) we can write:

$$(**) < \frac{\beta^{-1} + l + 2 \cdot d \cdot A(y_d)}{l - y_d - 2} \leq \frac{\beta^{-1} + l + 2 \cdot d \cdot A(y_d)}{\frac{l}{2}} \leq \frac{2 \cdot \beta^{-1}}{l} + 2 + \frac{4 \cdot d \cdot A(y_d)}{l} \leq \frac{1}{2} \cdot \beta^{-1} + 2 + d \cdot A(y_d)$$

where the last transition relies on the fact that $l \geq 4$.

Case II: $l \leq 2y_d + 3$. Notice that by definition l is always at least $y_d + 2$. Therefore, by continuing from Equation (2) we have:

$$(**) \leq \frac{\beta^{-1} + l + 2 \cdot d \cdot A(y_d)}{2 - 1 - \frac{A(y_d)}{d^{y_d}}} \leq \frac{\beta^{-1} + l + 2 \cdot d \cdot A(y_d)}{1 - \frac{1}{d-1}} = \frac{d-1}{d-2} (\beta^{-1} + l + 2 \cdot d \cdot A(y_d))$$

where for the second transition we used:

$$\frac{A(y_d)}{d^{y_d}} = \frac{d^{y_d} - 1}{(d-1)d^{y_d}} < \frac{1}{d-1}$$

Notice that $A(y_d) \geq y_d$, hence, $l \leq 2A(y_d) + 3$. Recall that $d \geq 3$ we can continue to bound (**):

$$(**) \leq \frac{d-1}{d-2} (\beta^{-1} + (2 \cdot d + 2) \cdot A(y_d) + 3) \leq 2 \cdot \beta^{-1} + 6 \cdot d \cdot A(y_d) + 6$$

From Case I, and case II combined, we have that:

$$(**) \leq 2\beta^{-1} + 6d \cdot A(y_d) + 6$$

□

Thus, if $k \geq 2\beta^{-1} + 6d \cdot A(y_d) + 6$, for any node v the strategy of choosing $c_i = l(v) - \varphi - 1$ for some child i is dominated by the strategy $(c_1, \dots, c_{i-1}, c_i - 1, c_{i+1}, \dots, c_d)$. Since k is the number of nodes outside the tree rooted by v then for the lemma to hold it is sufficient to have 7 seeds (each with d children that span trees at height y_d), and $2\beta^{-1} + 6$ additional nodes.

3 A Lower Bound on Dominant Strategy Schemes

In the previous sections, we presented several reward schemes and analyzed their behavior in equilibrium. The solution concept we analyzed was iterated removal of dominated strategies. Although iterated removal of dominated strategies is a strong solution concept, a natural goal is to seek good reward schemes that use even stronger solution concepts. We now show that there does not exist reward schemes that guarantee both low overhead and low payment, if the solution concept is dominant strategies.

Theorem 3.1 *Fix any $\mathcal{H}_s \leq H$ for each seed s . Let R be a reward scheme such that no duplication and information propagation is a dominant strategy equilibrium for every node up to depth \mathcal{H}_s that is in a tree rooted by a seed s . Suppose that at least half of the nodes are aware of the transaction at the end of the distribution phase. The expected payment of R is at least $\frac{1}{10} \left(\frac{2^{H-4}}{t^2} + \frac{1}{t} \cdot \left(\frac{H-3}{t \cdot e} \right)^{H-3} \right)$ where t is the number of initial seeds.*

As a corollary we get that for the total payment to be constant then the number of seeds must be $\Omega(2^{\frac{H}{2}})$. Moreover, for a constant t and large H the payment will be huge, at least $\Omega((H/c)^{H-3})$ for some constant c .

3.1 Proof of Theorem 3.1

Consider a reward scheme as in the statement. We claim that in a dominant strategy equilibrium, at least $\frac{1}{10}$ of the nodes that are aware of the transaction are in trees rooted by a seed s such that $\mathcal{H}_s \geq H - 2$. Otherwise, the number of nodes that are aware of the transaction is at most:

$$\frac{t}{10} \cdot \frac{d^H - 1}{d - 1} + \frac{9t}{10} \cdot \frac{d^{H-3} - 1}{d - 1} < \frac{t}{2} \cdot \frac{d^H - 1}{d - 1}$$

A contradiction to the assumption that at least half of the nodes are aware of the transaction. We will show that the expected payment (conditioned on a node in this tree hashing the transaction) of every tree with $H \geq \mathcal{H}_s \geq H - 2$ is at least $\frac{2^{\mathcal{H}_s-2}}{t^2} + \frac{1}{t} \cdot \left(\frac{\mathcal{H}_s-1}{t \cdot e} \right)^{\mathcal{H}_s-1} \geq \frac{2^{H-4}}{t^2} + \frac{1}{t} \cdot \left(\frac{H-3}{t \cdot e} \right)^{H-3}$, and the theorem will follow. From now on we fix such a tree and denote its payment simply by $r_{i,h}$ (dropping the superscript s). We start with several claims.

Claim 3.2 $r_{i,h} \geq r_{i,h+1} + r_{i+1,h+1}$.

Proof: Consider a node v and suppose that every child of v uses the strategy of cloning itself ($h-i-1$) times and does not distribute the message to its children. In case one of v 's children authorizes the transaction, the child declares that its last clone was the authorizer. Then, by Sybil proofness, v must not gain by duplicating itself another time and then sending the transaction to its children. \square

In the appendix we prove the following two claims. The first statement is similar to Pascal's triangle.

Claim 3.3 $r_{1,1} \geq \sum_{i=0}^{\mathcal{H}_s-1} \binom{\mathcal{H}_s-1}{i} \cdot r_{i+1,\mathcal{H}_s}$

Claim 3.4 *Let v be some node at position $h < \mathcal{H}_s$ that received the transaction. Let w be the number of nodes that are not descendants of v that are aware of the transaction. If v has a dominant strategy of transmitting the transaction without duplication then for every $1 \leq k \leq \mathcal{H}_s - h$, $\frac{r_{1,h}}{w} \leq r_{1+k,h+k}$.*

In the appendix we use Claim 3.2 and Claim 3.4 to prove the following lemma:

Lemma 3.5 *For any $m \geq 1$ it holds that: $r_{m,\mathcal{H}_s} \geq \frac{1}{2} \left(\frac{1}{t-1} + \frac{(m-1)!}{(t-1)^{m-1}} \right)$*

We are now ready to prove Theorem 3.1. Let t be the number of seeds. Suppose that $t - 1$ of them do not distribute the transaction, and consider the tree of the only seed that does distribute the transaction. By Lemma 3.5 it holds that for any $i \geq 1$ it holds that $r_{i, \mathcal{H}_s} \geq \frac{1}{2} \left(\frac{1}{t-1} + \frac{(i-1)!}{(t-1)^{i-1}} \right)$. Now, by Claim 3.3

$$\begin{aligned} r_{1,1} &\geq \sum_{i=0}^{\mathcal{H}_s-1} \binom{\mathcal{H}_s-1}{i} \cdot r_{i+1, \mathcal{H}_s} \geq \frac{1}{2} \sum_{i=0}^{\mathcal{H}_s-1} \binom{\mathcal{H}_s-1}{i} \cdot \left(\frac{1}{t-1} + \frac{i!}{(t-1)^i} \right) \geq \frac{1}{2} \left(\frac{2^{\mathcal{H}_s-1}}{t} + \frac{(\mathcal{H}_s-1)!}{t^{\mathcal{H}_s-1}} \right) \\ &\geq \frac{2^{\mathcal{H}_s-2}}{t} + \frac{\sqrt{2\pi(\mathcal{H}_s-1)}}{2} \cdot \left(\frac{\mathcal{H}_s-1}{t \cdot e} \right)^{\mathcal{H}_s-1} \geq \frac{2^{\mathcal{H}_s-2}}{t} + \left(\frac{\mathcal{H}_s-1}{t \cdot e} \right)^{\mathcal{H}_s-1} \end{aligned}$$

By applying Claim 3.4, we get that $\forall l \quad r_{h,h} \geq \frac{r_{1,1}}{t} \geq \frac{2^{\mathcal{H}_s-2}}{t^2} + \frac{1}{t} \cdot \left(\frac{\mathcal{H}_s-1}{t \cdot e} \right)^{\mathcal{H}_s-1}$

$r_{h,h}$ is the payment to the root in case the length of the winning chain was h , and so is a lower bound on the overall payment of the scheme in that case.

Acknowledgements

We thank Jon Kleinberg for helpful discussions and valuable comments.

References

- [1] The Bitcoin wiki. Available online at <https://bitcoin.it>.
- [2] The DARPA network challenge. Available online at <http://archive.darpa.mil/networkchallenge/>.
- [3] Joshua Davis. The crypto-currency: Bitcoin and its mysterious inventor. *The New Yorker*, October 10, 2011.
- [4] Kirill Dyagilev, Shie Mannor, and Elad Yom-Tov. Generative models for rapid information propagation. In *Proceedings of the First Workshop on Social Media Analytics, SOMA '10*, pages 35–43, New York, NY, USA, 2010. ACM.
- [5] Yuval Emek, Ron Karidi, Moshe Tennenholtz, and Aviv Zohar. Mechanisms for multi-level marketing. In *Proceedings of the 12th ACM conference on Electronic commerce, EC '11*, pages 209–218, New York, NY, USA, 2011. ACM.
- [6] Jon Kleinberg and Prabhakar Raghavan. Query incentive networks. In *FOCS '05: Proceedings of the 46th IEEE Symposium on Foundations of Computer Science*, pages 132–141, 2005.
- [7] Paul Krugman. Golden cyberfettters. Available online at <http://krugman.blogs.nytimes.com/2011/09/07/golden-cyberfettters/>.
- [8] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. Available online at <http://bitcoin.org/bitcoin.pdf>, 2008.
- [9] Galen Pickard, Iyad Rahwan, Wei Pan, Manuel Cebrián, Riley Crane, Anmol Madan, and Alex Pentland. Time critical social mobilization: The darpa network challenge winning strategy. *CoRR*, abs/1008.3172, 2010.
- [10] Fergal Reid and Martin Harrigan. An analysis of anonymity in the Bitcoin system. *CoRR*, arXiv/1107.4524, 2011.

- [11] James Surowiecki. Cryptocurrency. Technology Review, MIT. September/October, 2011. Available online: <http://www.technologyreview.com/computing/38392/>.

A A Formal Model

In our game, one transaction needs to be authorized. There are $n = t \cdot \frac{d^H - 1}{d - 1}$ nodes (players), organized in t complete d -ary trees each of height H , that are all directed towards the leaves. We call each root of a tree a *seed*. The *depth* of a node is the number of nodes on the path from the root of the node's tree to the node (the depth of the root is 1).

In the most general definition of the strategy space, a strategy of a node v is to specify for each $x \geq 1$ two things. First it specifies $\tilde{p}^{x,v} \geq 0$, the number of fake identities (duplicates) of itself v has added before trying to authorize the transaction with the last identity (it pretends that a node in a lower level succeeds in authorizing the transaction). Secondly, for every child i it specifies a pair $(b_i^{x,v}, \tilde{c}_i^{x,v})$. The meaning of the pair is that if v received the transaction from his parent and had $x - 1$ predecessors (including clones): $b_i^{x,v} \in \{0, 1\}$ specifies whether v sends the transaction to its i 'th child or not. In case v does send the transaction to its i -th child, $\tilde{c}_i^{x,v} \geq 0$ specifies how many fake identities of itself v has added before sending the transaction.

Given a strategy profile S , a player v is *aware of the transaction* if v is a seed, or if v is the i 'th child of u that has $x - 1$ predecessors, u is aware of the transaction, and pass it to v , that is $b_i^{x,u} = 1$.

Once a node authorizes the transaction it determines the *winning chain*, this is a sequence of identities on the path from a root node to the authorizing identity. Every reward scheme is completely defined by non-negative rewards $r_{i,h}^s$ for every seed s and every $1 \leq i \leq h$. The reward $r_{i,h}^s$ is the reward given to the i -th identity on the winning chain of length h that starts at authorizing node.

In this paper we consider reward schemes in which we want to motivate nodes not to duplicate themselves. Therefore, from now on we restrict our attention to schemes that never reward nodes with more than H predecessors, effectively implying that we only need to consider $x \leq H$. Hence, we can define a more concise representation of the strategy space that will be used throughout the paper.

Formally, the strategy S_v of a node v is represented as follows: for every $1 \leq l \leq H$, there is a tuple $S_v(l) = (c_1^{l,v}, \dots, c_d^{l,v})$ where $0 \leq c_i^{l,v} \leq l - 1$. In addition, for every $1 \leq l \leq H$ each node has a number $0 \leq p^{l,v} \leq l - 1$.

We say that a node v *propagates information and does not duplicate* at l in strategy profile S if $S_v(l) = (0, \dots, 0)$ and $p^{l,v} = 0$.

Given a strategy profile S , define the *level* $l(v)$ of a node v that is the i 'th child of u in the tree in a recursive way: if $l(u) = 0$ then $l(v) = 0$, and if $l(u) > 0$ then $l(v) = l(u) - 1 - c_i^{l(u),u}$ (note that $c_i^{l(u),u}$ is the i 's element of $S_u(l)$). The level of each seed is defined to be H . Informally, if v is aware of the transaction, then $H - l(v) + 1$ is the length of the chain from the seed to v (including clones).

Given a strategy profile S and a reward scheme we can define the utility of each node. An authorizer w is chosen uniformly at random among the players that are aware of the transaction and try to authorize it. Then, we allocate rewards to w and its ancestors. For each ancestor u of w let $q(u)$ be the number of duplicates of u in the path leading to w from the seed (that is for u which its i 'th child is in the path: $q(u) = c_i^{l(u),u}$). For w , let $q(w) = p^{l(w),w}$. Let $c(w) = H - l(w) + q(w) + 1$. The reward of any u in the path to w (including w) is $\sum_{i=0}^{q(u)} r_{c(w)-(i+H-l(u)+1),c(w)}^s$, where s is the seed that roots the subtree that w belongs to.

A reward scheme is *individually rational* for \mathcal{H} in a tree rooted by s if for every $1 \leq h \leq \mathcal{H}$ we have that $r_{1,h}^s \geq 1$. We assume that all players are expectation maximizers, so the utility of a player in a given profile is its expected utility, where expectation is taken over the selection of the authorizer w .

All logarithms in this paper have base d , so throughout the paper we write $\log H$ to denote $\log_d H$.

B A Brief Overview of Bitcoin

In this section we give a brief overview of the Bitcoin protocol. This is by no means a complete description. The main purpose of this section is to help understanding our modeling choices, and why our proposed reward schemes can be implemented within the context of the existing protocol. The reader is referred to [8, 1] for a complete description.

B.1 Signing Transactions and Public Key Cryptography

The basic setup of electronic transactions relies on public key cryptography. When Alice wants to transfer 50 coins to Bob, she signs a transaction using her private key. Hence, everyone can verify that Alice herself initiated this transaction (and not someone else). Bob, in turn, is identified as the target of the transfer using his public key. For the money to be actually transferred from Alice's account to Bob's account, some entity has to keep track of the lastest owner of the coins, and to change this owner from Alice to Bob. Otherwise, Alice could "double spend" her money – First transfer the coins to Bob, then transfer the same coins again to Charlie. Traditionally, this role was fulfilled by banks. In return, banks tended to charge high fees, for example in international transfers.

B.2 Agreeing on the History by Majority of CPU power

Bitcoin suggests a different solution to this problem. A peer to peer network is used to validate all transactions. Nodes in the network agree on a common history using a "majority of CPU power mechanism". A mechanism can not rely on a numerical majority of the nodes, as it is relatively easy to create additional identities in a network, for example by spoofing IP addresses. We first describe the protocol that the network implements and then explain why the history is accepted only if it is agreed upon by nodes that together control a majority of the CPU power.

Let us assume again that Alice wants to transfer 50 coins to Bob. Alice will send her signed transaction to some of the nodes in the network. Next, these nodes will forward the transaction to all of their neighbors in the network and so on. A node that receives the transaction first verifies that this is a valid transaction (e.g. that the money being transferred indeed belongs to Alice). If successful, the node adds this transaction to the block of transactions it attempts to authorize (the specific details on the structure of this block are omitted). To authorize a block, a node has to solve an inverse Hash problem. More specifically, its goal is to add some bits (nonce) to the block such that the Hash value of the new block begins with some predefined number of zeroes.

The number of zeroes is adjusted such that the average time it takes the network to sign a block is fixed. The protocol uses a clever method to aggregate transactions (a Merkle tree) which assures that the size of the block to be hashed is also fixed. Additional information that is included in the block is the hash of the previously authorized block. So, in fact, the authorized blocks form a chain in which each block identifies the one the preceded it.

When a node authorizes a block it broadcasts to the network the new block and the proof of work (the string which is added to the block to get the desired hash). If a node receives two different authorized blocks it adopts the one which is part of the longer chain.

Let us argue briefly and informally why the history is determined by the majority of the CPU power ([8]). That is, the probability that a group of malicious nodes manages to change the history decreases as the fraction of CPU they control decreases. Assume that a group of malicious nodes, that does not have the majority of the CPU power, wants to change the chain that has currently been adopted by the other nodes. Since the malicious nodes own less CPU power, their authorization rate is slower than the authorization rate of the majority of the network, which is honest. Therefore, the longer chains will be signed by the majority of the network (with high probability) and these are the ones that will be adopted by the honest nodes in the network. In fact, once a block has been accepted into the chain,

the probability that a longer chain that displaces it will appear decreases exponentially with time (as the chain that follows it grows longer it is harder to manufacture a longer alternative one). See [8] for a more formal argument.

B.3 Transaction Fees

To incentivize nodes to participate in the peer to peer network and invest effort in authorizing transactions, rewards have to be allocated. Currently, in the initial stages of the protocol this is done by giving the authorizer of a block a fixed amount of bitcoins (this also the only method for printing new bitcoins). This fixed amount will be reduced every few years at a rate determined by the protocol. As the money creation is slowly phased out, there will be a need to reward the nodes differently. The protocol has been designed with this in mind. It already allows the transaction initiator to specify a fee for authorizing her transaction. As we explained in the introduction this introduces an incentive problem since nodes prefer to keep the transaction to themselves instead of broadcasting it.

B.4 On Implementing the Hybrid Scheme

We start with a rough description of how fees for authorizing transactions in Bitcoin are currently implemented. Alice produces a transaction record in which she transfers some of her coins to Bob. She specifies the fee (if any) for authorizing this transaction in a field in the transaction record, and then cryptographically signs this record. When Victor authorizes a transaction, the implication is that Alice transferred some amount of money to Bob, and some amount of money (specified in the fee field) to Victor.

Any (β, \mathcal{H}) -almost uniform scheme can be implemented similarly (and thus, the Hybrid scheme can be implemented similarly). Alice produces a different transaction record for every seed, in which she specifies β , \mathcal{H} , and some amount of coins f (we normalized $f = 1$ in the previous sections – the total fee if a node in tree rooted by the seed authorizes the transaction will be $(1 + \beta \cdot \mathcal{H}) \cdot f$). We modify the protocol so that every node that transfers the transaction to its children cryptographically signs the transaction, and specifically identifies the child to whom the information is being sent using that child’s public key. The transaction information, therefore includes a “chain of custody” for the message which will be different for every node that tries to authorize the transaction. The implication of Victor authorizing a transaction is that Alice transferred some amount of money to Bob, a fee of $f \cdot \beta$ to every node in the path (as it appears in the authorized message), and a fee of $f + \beta \cdot (\mathcal{H} - h + 1) \cdot f$ to Victor if the path was of length $h \leq \mathcal{H}$. Payments will not be awarded if the “chain of custody” is not a valid chain that leads from Alice to Victor.

C Missing Proofs

C.1 Missing Proof of Lemma 2.6

Let us assume that some elimination order ends with a sub-game that does not contain any profile with full propagation and no duplication. Let T' be the last sub-game in the elimination order for which there is still a profile with full propagation and no duplication for every node at depth at most \mathcal{H} . It must be that for some player i in T' , the strategy s_i of full propagation and no duplication is dominated by another strategy s'_i in which this player either duplicates or does not fully distribute. In particular, let us fix the behavior of the other players to be the profile s_{-i} in which they fully distribute and do not duplicate (such a profile exists in T' by assumption). For s'_i to dominate s_i , it must be that $u_i(s_i, s_{-i}) \leq u_i(s'_i, s_{-i})$.

We will however show that the opposite holds, and thereby reach a contradiction. We define a series of strategies $s_i^1, s_i^2, s_i^3, \dots, s_i^m$ such that $s_i^1 = s'_i$, and $s_i^m = s_i$ for which we shall show:

$u_i(s'_i, s_{-i}) = u_i(s_i^1, s_{-i}) < u_i(s_i^2, s_{-i}) < \dots < u_i(s_i^m, s_{-i}) = u_i(s_i, s_{-i})$. Note, that the strategies $s_i^2, s_i^3, \dots, s_i^{m-1}$ may or may not be in T' , and that we are merely using them to establish the utility for s_i, s'_i .

The strategy s_i^{j+1} is simply the strategy s_i^j , with one change: player i replicates itself one less time to one of its children. Specifically, if player i replicates itself (c_1, \dots, c_d) times correspondingly for each of its d children, let $\nu = \argmax_j(c_j)$ it will instead replicate itself $(c_1, \dots, c_\nu - 1, \dots, c_d)$ times in the new strategy.

We establish the fact that $u_i(s_i^j, s_{-i}) < u_i(s_i^{j+1}, s_{-i})$ by a direct application of Claim 2.7. The claim shows that replicating one less time is better given sufficient external computation power, and given that the node's children fully distribute with no replication. Both of these are guaranteed in the profile s' .

C.2 Missing Proof of Claim 3.3

By induction on \mathcal{H}_s . The base case is trivial $r_{1,1} \geq r_{1,1}$. Next we assume that the claim holds for \mathcal{H}_s and show it also holds for \mathcal{H}_{s+1} . By the induction hypothesis, we have that: $r_{1,1} \geq \sum_{i=0}^{\mathcal{H}_s-1} \binom{\mathcal{H}_s-1}{i} \cdot r_{i+1, \mathcal{H}_s}$. We can now use Claim 3.2 to extend it to \mathcal{H}_{s+1} . We have that for every i : $r_{i, \mathcal{H}_s} \geq r_{i, \mathcal{H}_s+1} + r_{i+1, \mathcal{H}_s+1}$ and also $r_{i-1, \mathcal{H}_s} \geq r_{i-1, \mathcal{H}_s+1} + r_{i, \mathcal{H}_s+1}$. This implies that the coefficient of r_{i, \mathcal{H}_s+1} equals the sum of coefficients of $r_{i, \mathcal{H}_s} + r_{i-1, \mathcal{H}_s}$ which we have by the induction assumption. Hence, the coefficient of r_{i, \mathcal{H}_s+1} is $\binom{\mathcal{H}_s-1}{i-2} + \binom{\mathcal{H}_s-1}{i-1} = \binom{\mathcal{H}_s}{i-1}$ and the claim follows.

C.3 Missing Proof of Claim 3.4

If v does not distribute the transaction then it alone competes with w other nodes. If it does distribute to all of its children, the children use the following strategy: each child duplicates itself exactly $k-1$, and does not distribute the transaction further. In case a child u hashes the transaction then u declares itself as the end of a chain of length $h+k$: h nodes of v and k nodes that belong to u . In this case, the reward of v is $r_{1+k, h+k}$. Therefore, by the assumption that for every k , v prefers to distribute to all of his children even if they use the above strategies we have that: $\frac{r_{1,h}}{w} \leq \frac{r_{1,h} + d \cdot r_{1+k, h+k}}{w+d}$. By rearranging the previous inequality we have that $\frac{r_{1,h}}{w} \leq r_{1+k, h+k}$.

C.4 Missing Proof of Lemma 3.5

The proof is by induction on m . Recall that $r_{1, \mathcal{H}_s} \geq 1$ then the claim holds trivially for $m = 1$. We assume for smaller values of m and prove for $m + 1$. Our proof relies on the following claim:

Claim C.1 For any $h \leq \mathcal{H}_s$:

- $r_{\mathcal{H}_s-h+1, \mathcal{H}_s} \geq \frac{1}{t-1} \sum_{j=1}^{\mathcal{H}_s-h} r_{j, h+j}$
- $r_{1,h} \geq 1 + \frac{1}{t-1} \sum_{j=1}^{\mathcal{H}_s-h} r_{j, h+j}$

Proof: By Claim 3.4 for any k such that $h+k \leq \mathcal{H}_s$ it holds that

$$r_{1+k, h+k} \geq \frac{r_{1,h}}{t}.$$

which implies that

$$t \cdot r_{\mathcal{H}_s-h+1, \mathcal{H}_s} \geq r_{1,h}.$$

By Claim 3.2 for any $i \leq h < \mathcal{H}_s$ it holds that

$$r_{i,h} \geq r_{i+1, h+1} + r_{i, h+1}$$

which implies that

$$r_{1,h} \geq \sum_{j=1}^{\mathcal{H}_s-h} r_{j,h+j} + r_{\mathcal{H}_s-h+1,\mathcal{H}_s}$$

Now we can combine the above and derive that

$$t \cdot r_{\mathcal{H}_s-h+1,\mathcal{H}_s} \geq r_{1,h} \geq \sum_{j=1}^{\mathcal{H}_s-h} r_{j,h+j} + r_{\mathcal{H}_s-h+1,\mathcal{H}_s}$$

and thus

$$r_{\mathcal{H}_s-h+1,\mathcal{H}_s} \geq \frac{1}{t-1} \sum_{j=1}^{\mathcal{H}_s-h} r_{j,h+j}$$

and additionally by recalling that for every h , $r_{1,h} \geq 1$

$$r_{1,h} \geq 1 + r_{\mathcal{H}_s-h+1,\mathcal{H}_s} \geq 1 + \frac{1}{t-1} \sum_{j=1}^{\mathcal{H}_s-h} r_{j,h+j}$$

□

We are now ready to prove Lemma 3.5. By Claim C.1 we have that:

$$r_{m+1,\mathcal{H}_s} \geq \frac{1}{t-1} \sum_{j=1}^m r_{j,\mathcal{H}_s-m+j}$$

Notice that for all $0 \leq j \leq m$: $r_{j,\mathcal{H}_s-m+j} \geq r_{m,\mathcal{H}_s}$, by a simple generalization of the sybil proofness of Claim 3.2. We can now use the induction hypothesis and get that:

$$r_{m+1,\mathcal{H}_s} \geq \frac{1}{t-1} \left(r_{1,\mathcal{H}_s-m+1} + (m-1) \left(\frac{1}{2} \left(\frac{1}{t-1} + \frac{(m-1)!}{(t-1)^{m-1}} \right) \right) \right) \quad (3)$$

Also by using the second part of Claim C.1 we have that $r_{1,\mathcal{H}_s-m+1} \geq \left(1 + \frac{1}{t-1} \right) \sum_{j=1}^{m-1} r_{j,\mathcal{H}_s-m+1+j}$.

$$\begin{aligned} r_{1,\mathcal{H}_s-m+1} &\geq 1 + \frac{1}{t-1} (m-1) \left(\frac{1}{2} \left(\frac{1}{t-1} + \frac{(m-2)!}{(t-1)^{m-2}} \right) \right) \\ &\geq 1 + \left(\frac{1}{2} \left(\frac{(m-1)!}{(t-1)^{m-1}} \right) \right) \end{aligned} \quad (4)$$

By plugging in the bound from (4) in (3) we get that:

$$\begin{aligned} r_{m+1,\mathcal{H}_s} &\geq \frac{1}{t-1} \left(1 + \left(\frac{1}{2} \left(\frac{(m-1)!}{(t-1)^{m-1}} \right) \right) + (m-1) \left(\frac{1}{2} \left(\frac{(m-1)!}{(t-1)^{m-1}} \right) \right) \right) \\ &\geq \frac{1}{t-1} \left(1 + \frac{1}{2} \left(\frac{m!}{(t-1)^{m-1}} \right) \right) \\ &\geq \frac{1}{2} \left(\frac{1}{t-1} + \frac{m!}{(t-1)^m} \right) \end{aligned}$$